Moore's Law

The number of transistors on a chip has doubled approximately every 12–18 months for the past 30 years—a trend referred to as Moore's Law. (See figure 9-2.) This trend is named for Gordon Moore of Intel, who first observed it. As Moore (1999) noted:

> I first observed the "doubling of transistor density on a manufactured die every year" in 1965, just four years after the first planar integrated circuit was discovered. The press called this "Moore's Law" and the name has stuck. To be honest, I did not expect this law to still be true some 30 years later, but I am now confident that it will be true for another 20 years.

Performance has increased along with the number of transistors per chips, while the cost of chips has remained fairly stable. These factors have driven enormous improvements in the performance/cost ratio. (See figure 9-3.)

The complexity and cost of developing new chips and new semiconductor manufacturing equipment also have increased. As a result, the industry has been driven toward greater economies of scale and industry-wide collaboration. Moore's Law—which began as the observation of an individual in a single company—has become a self-fulfilling prediction that drives industry-wide planning. Since 1992, the U.S. Semiconductor Industry Association (SIA) has developed a National Technology Roadmap for Semiconductors, which charts the steps the industry must take to maintain its rate of improvement. In 1998, this effort evolved into the International Technology Roadmap for Semiconductors, with participation by the Japanese, European, and South Korean semiconductor industries. The 1998 update projects the number of transistors per chip increasing to 3.6 billion in 2014 (SIA 1998).

Figure 9-2.
Moore's Law

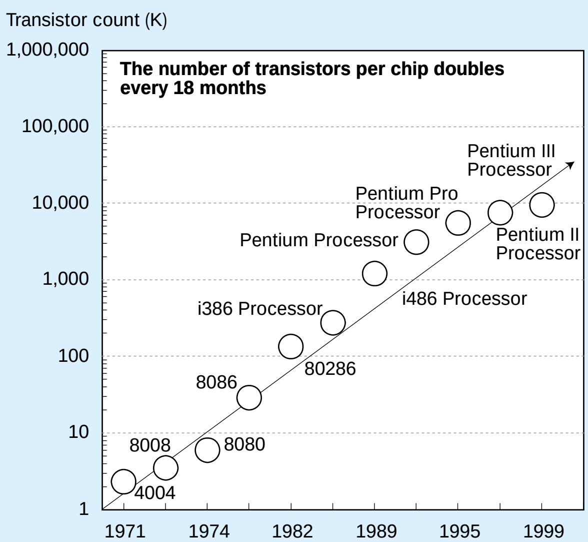

SOURCE: Intel. Available from <<http://www.intel.com/pressroom/kits/processors/quickref.htm>>.

See appendix table 9-1. *Science & Engineering Indicators – 2000*

Figure 9-3.
Price index for memory chips and microprocessors

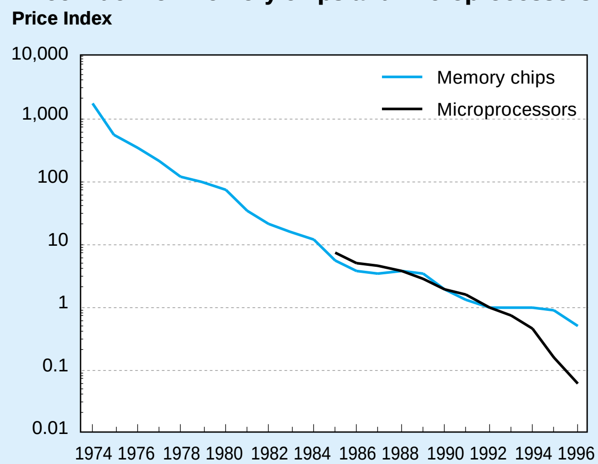

NOTE: 1992 = 100 (Log scale)

SOURCE: Grimm, B.T. "Price Indexes for Selected Semiconductors, 1974–96." *Survey of Current Business* (February 1998). Available from <<http://www.bea.doc.gov/bea/ARTICLES/NATIONAL/NIPA/1998/0298od.pdf>>.

Science & Engineering Indicators – 2000

more than 56 million Internet hosts—computers connected to the Internet—in July 1999, up from about 30 million at the beginning of 1998. (See figure 9-4 and appendix table 9-2.)

Information Technology Over the Past 50 Years

IT and the National Science Foundation (NSF) have come of age together. In this year that marks the 50th anniversary of NSF, few areas demonstrate as vividly as IT the progress that has been made in science and engineering in the past half-century.

In 1945, the same year that Vannevar Bush outlined his ideas for what became the National Science Foundation in *Science—the Endless Frontier*, he also wrote an article in the *Atlantic Monthly* that described his vision for capturing and accessing information. (See sidebar, "Excerpts from 'As We May Think'.") In the *Atlantic* article, Bush proposed the development of a kind of work station, which he called a "memex," that would store and provide access to the equivalent of a million volumes of books. The memex would also employ a way of linking documents "whereby any item may be caused at will to select immediately and automatically another"—allowing the user to build a trail between multiple

Figure 9-4.
Internet domain survey host count

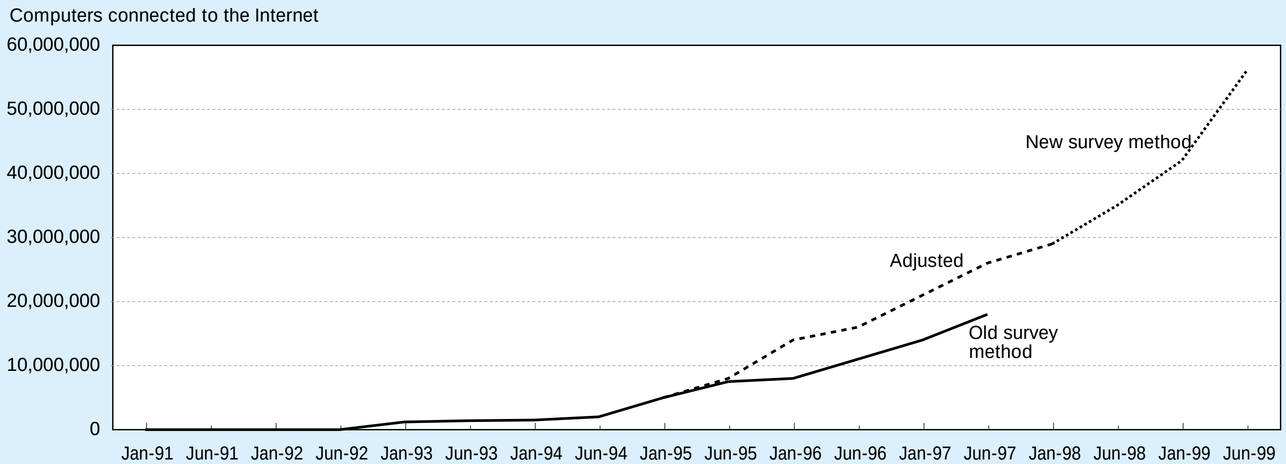

SOURCE: Internet Software Consortium. Available from <<http://www.isc.org/>>.

See appendix table 9-2.

Science & Engineering Indicators – 2000

documents. Although Bush proposed using photographic methods for storage and mechanical means for retrieval, and the exact technological capability he dreamed of has not yet come to pass, the proposed function of his memex is remarkably similar to hypertext today.

When Bush thought about the capabilities that would be dramatically useful to knowledge workers, he envisioned not capable calculators or word processors but capabilities to store and access information that current technology is just now achieving—using quite different approaches. Much R&D and innovation have been necessary to reach these capabilities.

In the same year that Bush's *Atlantic* article appeared, developments were taking place that would provide a different path for achieving his vision. At the University of Pennsylvania, John P. Eckert and John W. Mauchly were completing, with Army funding, what is commonly recognized as the first successful high-speed digital computer—the ENIAC. Dedicated in January 1946 and built at a cost of $487,802 (Moye 1996), the ENIAC used 18,000 vacuum tubes, covered 1,800 square feet of floor space, and consumed 180,000 watts of electrical power. It was programmed by wiring cable connections and setting 3,000 switches. It could perform 5,000 operations per second (CSTB 1998).

Also in 1945, Hungarian-born Princeton mathematician John von Neumann developed the stored program concept, which enabled computers to be programmed without rewiring. The von Neumann architecture—which refers to a computer with a central processing unit that executes instructions sequentially; a slow-to-access storage area; and secondary fast-access memory—became the basis for most of the computers that followed. Since the middle of the 20th century, software development has emerged as a discipline with its own challenges and skill requirements, complementing the more visible advances in hardware and enabling great systems complexity.

Over the succeeding 50 years, a vast number of innovations and developments occurred. (See sidebar, "IT Timeline.")

Innovations in IT over this period came from a remarkable diversity of sources and institutional settings, as well as a remarkable interplay among industry, universities, and government. Transistors and integrated circuits were invented by industry. Early computers and advances such as core memory, time-sharing, artificial intelligence, and Internet browsers were developed in universities, primarily with government funding. The World Wide Web was developed at the European Center for Particle Research (CERN), a high-energy physics laboratory. The mouse and windows were developed at a nonprofit research institute, with government funding. High-performance computers were mostly developed in industry with federal funds and with the involvement of federal laboratories. The diversity and close interaction between these institutions clearly contribute to the vitality of innovation in IT.

Innovation in IT has benefited from the support of a diverse set of federal agencies—including the Department of Defense (DOD), including the Defense Advanced Research Projects Agency (DARPA) and the services; NSF; the National Aeronautics and Space Administration (NASA); the Department of Energy (DOE); and the National Institutes of Health (NIH). Federal support has been particularly important in long-range fundamental research in areas such as computer architecture, computer graphics, and artificial intelligence, as well as in the development or procurement of large systems that advanced the technology—such as ARPANET, the Internet (See sidebar "Growth of the Internet"), and high-performance computers (CSTB 1998).

Often there has been complementary work supported by the Federal Government and industry. In many cases the Federal Government has supported the initial work in technolo-

Excerpts from "As We May Think"

Atlantic Monthly (July 1945)
by Vannevar Bush

Professionally our methods of transmitting and reviewing the results of research are generations old and by now are totally inadequate for their purpose...The difficulty seems to be, not so much that we publish unduly in view of the extent and variety of present day interests, but rather that publication has been extended far beyond our present ability to make real use of the record. The summation of human experience is being expanded at a prodigious rate, and the means we use for threading through the consequent maze to the momentarily important item is the same as was used in the days of square-rigged ships.

Consider a future device for individual use, which is a sort of mechanized private file and library. It needs a name, and, to coin one at random, "memex" will do. A memex is a device in which an individual stores all his books, records, and communications, and which is mechanized so that it may be consulted with exceeding speed and flexibility. It is an enlarged intimate supplement to his memory.

It consists of a desk, and while it can presumably be operated from a distance, it is primarily the piece of furniture at which he works. On the top are slanting translucent screens, on which material can be projected for convenient reading. There is a keyboard, and sets of buttons and levers. Otherwise it looks like an ordinary desk.

In one end is the stored material. The matter of bulk is well taken care of by improved microfilm. Only a small part of the interior of the memex is devoted to storage, the rest to mechanism. Yet if the user inserted 5,000 pages of material a day it would take him hundreds of years to fill the repository, so he can be profligate and enter material freely. It affords an immediate step…to associative indexing, the basic idea of which is a provision whereby any item may be caused at will to select immediately and automatically another. This is the essential feature of the memex. The process of tying two items together is the important thing.

When the user is building a trail, he names it, inserts the name in a code book, and taps it out on his keyboard. Before him are the two items to be joined, projected onto adjacent viewing positions. At the bottom of each there are a number of blank code spaces, and a pointer is set to indicate one of these on each item. The user taps a single key, and the items are permanently joined. In each code space appears the code word. Out of view, but also in the code space, is inserted a set of dots for photocell viewing; and on each item these dots by their positions designate the index number of the other item.

Thereafter, at any time, when one of these items is in view, the other can be instantly recalled merely by tapping a button below the corresponding code space. Moreover, when numerous items have been thus joined together to form a trail, they can be reviewed in turn, rapidly or slowly, by deflecting a lever like that used for turning the pages of a book. It is exactly as though the physical items had been gathered together from widely separated sources and bound together to form a new book. It is more than this, for any item can be joined into numerous trails.

The owner of the memex, let us say, is interested in the origin and properties of the bow and arrow. Specifically he is studying why the short Turkish bow was apparently superior to the English long bow in the skirmishes of the Crusades. He has dozens of possibly pertinent books and articles in his memex.

First he runs through an encyclopedia, finds an interesting but sketchy article, leaves it projected. Next, in a history, he finds another pertinent item, and ties the two together. Thus he goes, building a trail of many items. Occasionally he inserts a comment of his own, either linking it into the main trail or joining it by a side trail to a particular item. When it becomes evident that the elastic properties of available materials had a great deal to do with the bow, he branches off on a side trail which takes him through textbooks on elasticity and tables of physical constants. He inserts a page of longhand analysis of his own. Thus he builds a trail of his interest through the maze of materials available to him.

gies that were later developed by the private sector. In other cases Federal research expanded on earlier industrial research. Higher-level computer languages were developed in industry and moved to universities. IBM pioneered relational databases and reduced-instruction-set computing, which were further developed with NSF support. Collaboration between industry and university researchers has facilitated the commercialization of computing research. (See figure 9-5.)[1]

Most of the relentless cost-cutting that has been so important in the expansion of IT has been driven by the private sector in response to competitive pressures in commercial markets, although here too federal investment—such as in semiconductor manufacturing technologies—has played an important role in some areas.

[1]For a more complete description of industry and government roles in developing information technologies, see CSTB (1998).

IT Timeline

1945: The ENIAC, the first high-speed digital computer, is built at the University of Pennsylvania for the U.S. Army's Ballistics Research Laboratory to help prepare artillery firing tables.

1947: Bell Telephone Laboratories develops the transistor.

1949: The concept for core memory is patented by An Wang at Harvard University. Core memory and random access memory (RAM) are further developed by the Whirlwind Project at MIT.

1951: UNIVAC, the first commercial computer, is developed and delivered to the Census Bureau.

1952: G.W. Dummer, a radar expert from the British Royal Radar Establishment, proposes that electronic equipment be manufactured as a solid block with no connecting wires; he receives little support for his research.

1953: IBM enters the computer business with the 700 series computer.

1959: Texas Instruments and Fairchild Semiconductor both announce the integrated circuit.

Late 1950s–early 1960s: Timesharing (the concept of linking a large numbers of users to a single computer via remote terminals) is developed at MIT.

1961: Fairchild Semiconductor markets the first commercial integrated circuits.

1964: The IBM 360 is introduced and becomes the standard institutional mainframe computer.

1965: Gordon Moore predicts that the number of components in an integrated circuit will double every year (Moore's Law).

1968: Doug Engelbart of Stanford Research Institute demonstrates a word processor, a mouse, an early hypertext system, and windows. Gordon Moore and Robert Noyce found Intel.

1969: ARPANET goes online. Xerox establishes the Palo Alto Research Center to explore the "architecture of information."

1970: Fairchild Semiconductor introduces a 256-bit RAM chip.

1971: Intel introduces the 4004, a 4-bit microprocessor.

1972: Intel introduces the 8008, the first 8-bit microprocessor. E-mail is introduced over ARPANET.

1973: Robert Kahn and Vinton Cerf develop the basic ideas of the Internet.

1975: The MITS Altair 8800 is hailed as the first "personal" computer. Paul Allen and Bill Gates develop BASIC for the Altair 8800.

1976: Microsoft and Apple are founded.

1977: Apple markets the Apple II for $1,195; it includes 16K of RAM but no monitor.

1979: Software Arts develops the first spreadsheet program, Visicalc, which is an immediate success.

1981: The IBM PC is released.

1982: TCP/IP (Transmission Control Protocol and Internet Protocol) is established as a standard for ARPANET.

1984: The Apple Macintosh is released, featuring a simple, graphical interface.

1986: NSF establishes NSFNET and five supercomputing centers.

1987: The number of network hosts exceeds 10,000.

1989: The number of network hosts exceeds 100,000.

1989: Microsoft's annual sales reach $1 billion. The World Wide Web is developed at CERN.

1992: The number of Internet hosts exceeds 1 million.

1993: Mosaic, the first Web browser, is developed at the NSF-funded National Center for Supercomputer Applications at the University of Illinois, leading to rapid growth of the World Wide Web.

1994: Main U.S. Internet backbone traffic begins routing through commercial providers.

1995: NSFNET privatized.

SOURCES: PBS Online companion Web site for television special "Triumph of the Nerds: The Rise of Accidental Empires," <<http://www.pbs.org/nerds/timeline/micro.html>>; Virginia Tech Virtual Museum of Computing, Chronology of Events in Computer History, <<http://video.cs.vt.edu:90/cgi-bin/Lobby?Method=Chronology>>; Leiner et al. (1998).

Figure 9-5.
Government support for computing research

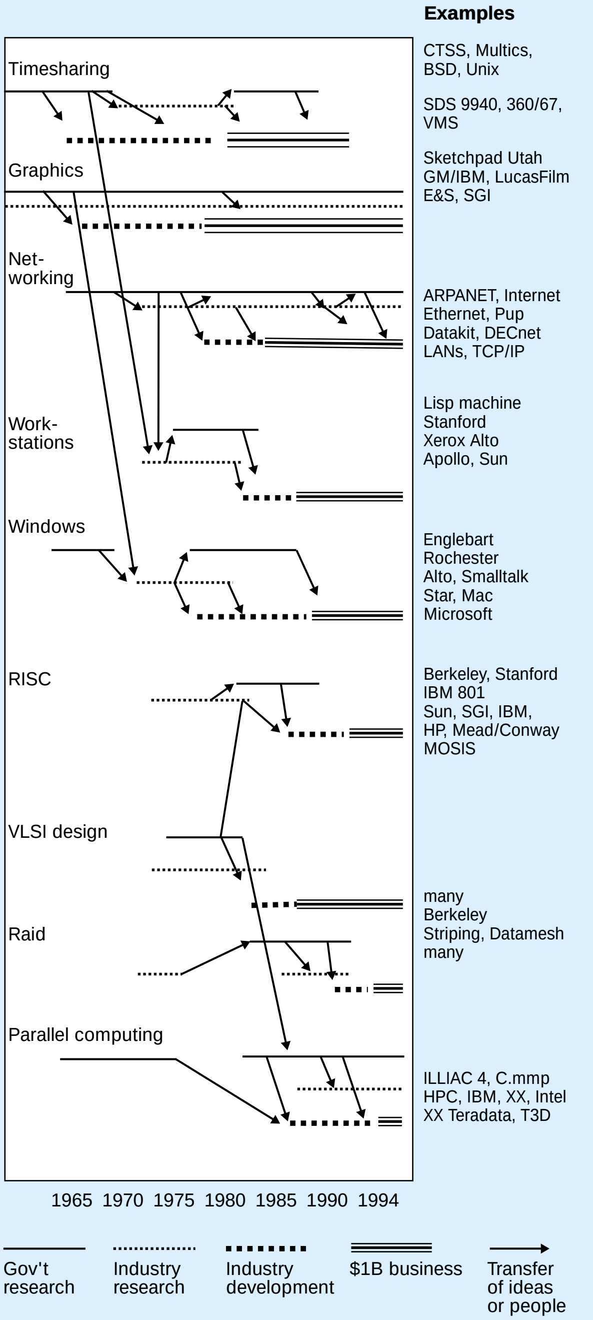

SOURCE: National Research Council, Computer Science and Telecommunications Board, *Funding a Revolution; Government Support for Computed Research* (Washington, DC: National Academy Press, 1999). *Science & Engineering Indicators – 2000*

Growth of the Internet

The Internet is a meta-network for a variety of subnetworks and applications such as the World Wide Web, bulletin boards, Usenet newsgroups, e-mail, scientific data exchange, and more. The foundation for the Internet was ARPANET, which started as four computer nodes in 1969. ARPANET was initiated by DARPA and was based on a then-new telecommunications technology called packet switching. ARPANET flourished as a medium for information and data exchange among universities and research laboratories. Moreover, it stimulated the development of TCP/IP, a communications protocol that enabled the interconnection of diverse networks. By the late 1970s, ARPANET comprised hundreds of computer nodes and integrated several separate computer networks, including one based on satellite technology.

The Internet grew out of the ARPANET, which converted to the TCP/IP protocol in 1983. NSF sponsored CSNET and later NSFNET (a high-speed network to link supercomputing centers), which became the backbone for the Internet. NSFNET replaced ARPANET in 1990 and expanded to include a variety of regional networks that linked universities into the backbone network. Many smaller networks linked into NSFNET. By early 1994, commercial networks became widespread, and almost half of all registered users of the network were commercial entities.

Two other events dramatically reshaped the character of the Internet. First, in 1989, scientists at CERN developed the World Wide Web and introduced it in experimental form. Second, in 1993, a team of programmers at NSF's National Center for Supercomputing Applications at the University of Illinois introduced Mosaic, a graphical (hypermedia) browser for exploring the Web. Mosaic was made available on the Internet at no cost, and Web use soared. (See figure 9-4.)

NSFNET was fully privatized in 1995, when there were enough commercial Internet service providers, Web browsers, and search engines to sustain the network's operations and management. The Internet continues to evolve. The Next Generation Internet Initiative is developing a higher-speed, more functional telecommunications network.

For more information on the Internet, see Cerf (1997) and Leiner et al. (1998).